Towards a Crowdsourced Radio Map for Indoor Positioning System

Ran Guan
Computer Laboratory
University of Cambridge
Cambridge, UK
Email: ran.guan@cl.cam.ac.uk

Robert Harle
Computer Laboratory
University of Cambridge
Cambridge, UK
Email: robert.harle@cl.cam.ac.uk

Abstract—Fingerprinting has become the most popular approach for infrastructure-free indoor positioning systems. But fingerprinting relies on frequent and exhaustive building surveys to build and maintain a radio map of the indoor environment. A concept being developed recently is the crowdsourced indoor positioning system, where users of an indoor area will collectively contribute sensor data collected by personal smart devices to the construction of the radio map. As a proof of concept, this paper proposes and evaluates a promising approach for building the radio map based on a crowdsourced dataset. Evaluation shows that we can achieve a mean positioning error of less than 3 meters base on the crowdsourced map, slightly worse than conventional path survey approach.

I. INTRODUCTION

A variety of infrastructure-based or opportunistic Indoor Positioning Systems (IPS) have been proposed and explored over the last two decades. Due to the ubiquity of 2.4 GHz radio signals like WiFi and Bluetooth, fingerprinting based on these opportunistic signals has become a well-established and mainstream approach. Typical fingerprinting systems have two stages. Firstly, an offline surveying stage, during which a dedicated surveyor travels extensively through the target area and builds a radio map from the collected fingerprints. Secondly, an online positioning stage, during which a building user observes a new fingerprint and compares it to the radio map to estimate location. At present, the accuracy of a fingerprinting IPS ranges from 2 to 10 m in laboratory conditions [1]. Although fingerprinting requires no extra infrastructure, it is notorious for being costly in terms of deployment and maintenance effort. According to [2], a 450,000 m^2 site survey took 15 surveyors two weeks to finish. The expensive construction of the radio map deters wide adoption of fingerprinting positioning systems, and motivates research into how to simplify and automate the survey [3].

Crowdsourcing is a promising alternative to the dedicated surveying process: people move around their environment carrying smart devices, which could be used to opportunistically survey it. However, the major challenge for this approach is that, while devices can supply detailed signal measurements, they rarely know where they are, preventing the fingerprints from assigned to a location.

In this work we describe a technique that seeks to align *sequences* of signal measurements collected as users walk through their environment. These sequences derive from the user taking short walks that we can characterise using Pedestrian Dead Reckoning (PDR) techniques. We align them in an arbitrary spatial frame using the signals they are trying to map and eventually embed the aligned sequences into the building's frame of reference. This provides enough information to create radio maps of the environment.

We test the technique using Bluetooth Low Energy (BLE) signals. The choice of BLE over the WiFi is due to the increasingly slow WiFi scan rate of consumer mobile devices. As the WiFi specification has grown to use other parts of the spectrum (e.g. the 5 GHz ISM band), the duration of a scan has also grown. A single scan can take 3–5 seconds on some consumer hardware. In this time a typical pedestrian will move approximately 5 m, making it difficult to associate the scan results with a meaningful position. BLE beacons can be configured with a wide range of update rates, allowing more detailed evaluation.

In this paper we also assume that the environment is corridor-based. This allows us to develop the core algorithms and helps with intuition. The rest of the paper is organised as follows. Section two briefly introduces related work. Section three describes the alignment and map generation methods. Section four details our experimental validation, which used 10 participants contributing collectively 100 walks. Section five concludes.

II. RELATED WORK

A signal fingerprint usually includes the received signal strength (RSS) values from multiple access points (APs) and an associated location tag. If sufficient fingerprints are collected to cover the entire target area, a Gaussian Process or similar regression approach can be applied to construct a 2D radio map. Although collecting RSS is straightforward with modern personal smart devices, recovery of the location tag is challenging even to dedicated surveying techniques. Conventionally, the surveyor has to either stand on predefined gridpoints or walk along predefined paths collecting fingerprints so that the location tag can be easily retrieved. In a crowdsourced scenario, getting the location tag is more difficult as users are usually not constrained and not aware of their accurate position.

In the early exploration of a crowdsourced IPS like Red-pin [4] and [5] in 2010, a fingerprint was generated when a user explicitly tagged a location to associated the ground-truth with the current signal observation. But in essence, these are little different from conventional surveying methods except that they require users to play the role of dedicated surveyors. They suffer from many of the same problems.

MIT introduced the concept of an organic location system [6], which still relied on user-generated location tags. But the organic system outlined a more detailed analysis of how the radio map grows in a piecemeal fashion in the crowdsourced context and how the uncertainty of the radio map varies spatially.

With the recent advances in microelectromechanical systems, inertial measurement units (IMU) on modern smart phones became accurate enough to allow various Pedestrian Dead Reckoning (PDR) algorithms to estimate a walking trajectory of a user. Based on the estimated walking trajectory, Zee [7] used an augmented particle filter and floor plan to estimate the location tag of a crowdsourced fingerprint. The particle filter not only modelled the position of the user, but also the step length and the heading of the user to fine-tune the PDR algorithm. However, in this scenario the user may have to walk for some distance before the particle filter can lock onto a unique path, which makes Zee less realistic for crowdsourcing.

Yang et al. proposed representing all fingerprints in graph nodes [8]. The number of steps taken to connect two locations are also taken into account while building the radio map. Consequently, all data can be treated in an unsupervised way, where rooms and corridors will be clustered into several groups in the signal space. Location tags can be therefore determined by matching to these clusters and the entire process assumes no a-priori information such as the floor plan or positions of the access points (APs).

All of these previous works view crowdsourcing as a two-phase process: firstly, attach location tag(s) to each fingerprint (or sequence of fingerprints) separately; secondly, combine the fingerprints intelligently to form a global map. The latter step is complicated by the noisy characteristics of crowdsourced data. The approach we take here considers these phases as a single optimisation phase.

III. CROWDSOURCING THE RADIO MAP

As discussed above, an effective radio map is built from fingerprints with reliable location tags. The system we are proposing does not try to location-tag individual fingerprints from the beginning but starts from aligning and merging different crowdsourced sequences, which are then tagged as a whole. The entire system can be roughly divided into five parts:

1) data crowdsourcing;
2) walk identification and characterisation (PDR);
3) approximate pair-wise alignment of sequences in signal space;
4) global alignment optimisation and map generation; and

5) embedding the map in the building frame of reference.

In this paper we focus primarily on tasks 2–4. Task 1 (data crowdsourcing) is essentially an engineering task, although it may be complicated by issues such as anonymity, which are beyond the scope of this work. Embedding into the building frame of reference is also straghtfowrad if we assume that task 4 suceeded and a floorplan is available. We now focus on tasks 2–4 individually.

A. Walk Identification and Characterisation

Walk and step detection is now widely available as a feature of mobile platforms such as Google Android. We form a PDR system with the inclusion of gyroscope measurements to output trajectory estimates formed by concatenating step vectors [9], [10]. These trajectories are, however, not embedded in any meaningful frame of reference, and may be scaled due to a lack of information on step length (in this paper we use a set step length of 70 cm in the first pass).

PDR is inherently unreliable over long timescales due to sensor drift and the possibility of outlier events (a stumble for example). However we have two advantages here: in-building walks are typically short; and we expect to have many walks collected by multiple people, allowing us to discard everything with an even remotely suspicious event.

B. Loop Closure Detection and Alignment

Given two separate radio sequences and associated PDR trajectories, loop closure detection aims to match points of the trajectories that have very similar signal properties (suggestive of them being at the same position and hence anchoring the two trajectories relative to each other). We wish to establish the translation, rotation and scaling factor between each pair of overlapping trajectories.

Two trajectories do not overlap simply because they look similar, so more information from the signal space is necessary to confirm a matching. In robotics, particularly visual Simultaneous Localisation and Mapping (SLAM), loop closures are often determined by detecting the same visual features. For fingerprinting, magnetic sequences can be utilised as an alternative to visual clues if the phone is guaranteed to be in the same position/orientation and entire the survey occurs over a short period of time [3]. In a crowdsourcing scenario, however, loop closures from one user are unlikely in one sequence and instead we combine multiple sequences from potentially different days and users.

Our current alignment approach is optimised for buildings with corridors (i.e. typical workplaces). We segment each PDR trajectory into sections punctuated by significant turns and search for (near) straight-line segments of 10 m or more (strong candidates for being corridors). It is these corridor segments that we then try to match. If we are able to align and embed the corridor segments, other segments can then be embedded through their relative relationship with the corridor segments.
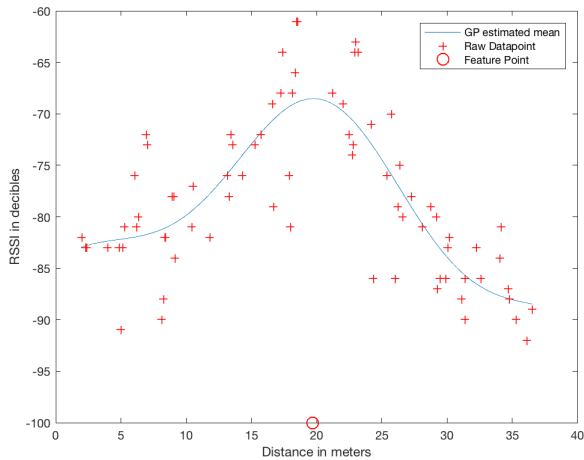
Fig. 1: 1-D GP of a BLE beacon

Loop closure detection and alignment   All merged fingerprints for one BLE

Fig. 2: Partial Results

1) Approximate Alignment of Two Segments: This step takes pairs of segments and aligns them by comparing the signals strengths observed along them. The output is one or more sets of segments that are believed to share part of their trajectory, and approximate (relative) alignments. The alignment is then globally optimised in the next step.

A given segment consists of a PDR trajectory estimate and a sequence of signal fingerprints collected while traversing that trajectory. Here a fingerprint is a set of signal strengths associated with different signal sources observed in a single scan (i.e. approximately at the same moment). Our initial attempts to match two signal sequences involved point-to-point mapping based on comparing individual fingerprints from on segment to those of the other. Our assumption was that small fingerprint distances (we used Euclidean distance in signal space) indicated likely co-location. Unfortunately the large noise on typical indoor signal strength measurements meant that this point-wise approach did not work well in practice.

Instead we use a map-to-map matching approach. We use Gaussian Process (GP) regression [11] to generate signal strength maps from each segment individually and then align the GP *maps* rather than the raw fingerprints. Figure 1 illustrates how the GP regression extracts the expected trend in signal strength when walking past the source.

Given a pair of segments suspected to overlap, we look for signal sources that produced a strong maximum in signal strength (as per Figure 1). We then compute the transformations that would align these maxima best.

In practice, this pairwise alignment of segments gives only approximate alignment. The radio measurements are very noisy and relatively sparse so any regression will not be able to pinpoint the true maximum position accurately. However, the alignment is typically close enough to initialise a global optimisation of alignments.
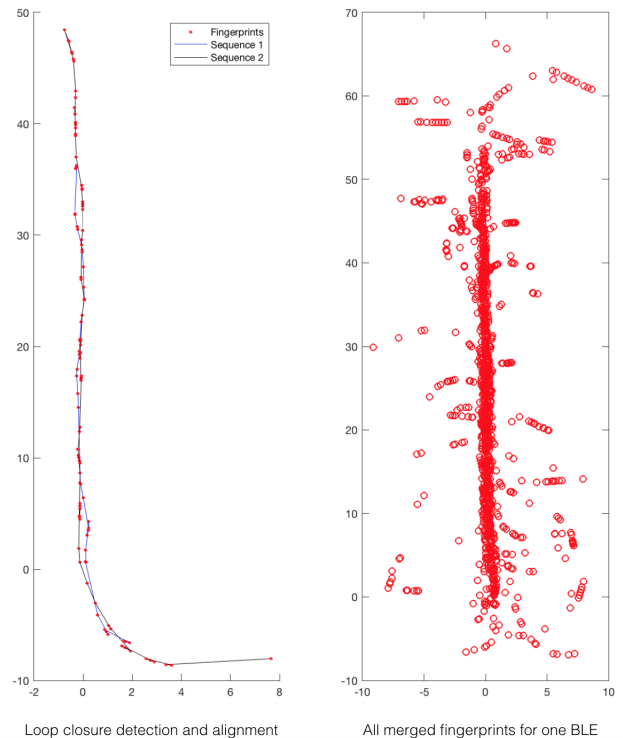
C. Global Optimisation

This step optimises the approximate alignments from the previous step. The process is inspired by GP regression itself, which optimises the log marginal likelihood (LML [12]) of observing the training dataset given the hyperparameters, which is a good measure of the goodness of fit of the output map to the input data

Intuitively, the quality of the sequence alignment will also be reflected in this goodness of fit measure: incorrect alignment will give spurious locations for measurements and cause the map variance to increase in those areas. Therefore we optimise the LML with regards to the alignment parameters (i.e. position, rotation and scaling that align one sequence to another). Clearly the state space is very large, making batch learning infeasible. Additionally, the closed form of the LML with respect to the alignment is not clear yet, meaning any optimisation approach with closed form derivatives is also infeasible. Consequently we start by optimising the longest sequences using the simplex search method [13] (with the heuristic alignment from the previous stage as the initial simplex) so their alignments can serve as optimisation constrains for the shorter sequences. At the end of the optimisation, we have generated signal maps in a single frame of reference. However, this frame of reference has not been anchored to a global (building) reference.

D. Embedding the Map

To embed the map in the building's reference frame in the general case is out of scope for this work. To allow
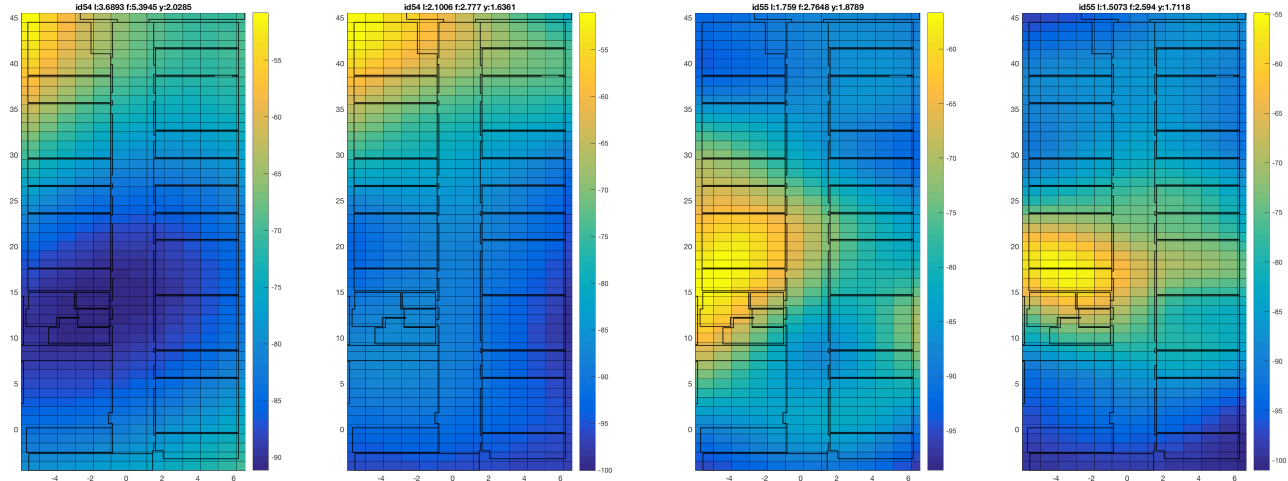
| (a) Crowdsourced on BLE A | (b) Surveyed on BLE A | (c) Crowdsourced on BLE B | (d) Surveyed on BLE B |

Fig. 3: Crowdsourced and Surveyed Radio Maps

quantitative assessment, we assume that the map is being generated for a building structured around corridors. In this case the estimated relative positions from the previous step (see Figure 2 for example) gives a characteristic shape. As our testbed is a rectangular corridor, we use a Principal Component Analysis (PCA) to identify the major axis of these positions, with the first component representing the corridor direction (highest variation on the corridor direction), and we therefore create a rectangular bounding box with longer edges parallel to the first component and shorter edges perpendicular to the first component. We perform a non-reflective similarity transformation to map the bounding box to our corridor's floor plan. This allows every fingerprint to be location-tagged onto the floor plan accordingly and a GP can be applied to construct the radio map afterwards.

IV. EVALUATION

A. Data Gathering

To evaluate our approach, we invited ten participants (eight males and two females) to simulate a crowdsourced scenario in a corridor in our building over the course of one day. We chose to simulate the walks to allow us to have a ground truth of the movements; to ensure good spatial coverage over a short period; and to ensure good inertial inputs for PDR (the PDR algorithms were not the focus of this work).

For each participant, a unique room was assigned as a nominal 'workplace', and also five other destinations: two ends of the corridor and three randomly chosen rooms along the corridor. The participant was asked to travel between these five destinations and the workplace under observation. Each participant held a Google Nexus 5X smart phone, which logged the inertial data and signal strengths from 12 Bluetooth Low Energy (BLE) beacons installed in the environment.

Participants usually took less than 10 minutes to finish all walks and room-level ground truth of these sequences was recorded for reference. The corridor was in use by other people during the experiment and there were typical levels of human activity throughout.

Major Sequence Length	Number of Sequences
0~10m	58
10~20m	24
20~30m	12
30~40m	5
>40m	1

TABLE I: Length of Longest Straight Line

The distribution of the length of the longest straight line of the recorded sequences is given in Table I. A significant number of walks were short walks. For these, the crowdsourced sequence proved too limited to be merged into the larger fingerprint library. In fact, when the length was less than 12 meters, identifying more than two BLE feature points to perform the initial match on was problematic. Ultimately we used the longest 25 sequences as input to our system. This gave a fingerprint library with around one thousand fingerprints for each BLE beacon. Figure 2 illustrates the inputs fingerprint locations.

After the experiment, we collected two further datasets. Firstly, we conducted a manual survey covering all the areas visited during the experiment, which took the surveyor around 30 minutes to finish. We implemented the approach from [14] on a Google Tango tablet as the ground-truth surveying tool for a manual survey and sampled signal observations using the same type of Nexus 5X used in the simulated crowdsourcing to avoid hardware differences in measurement. We build another radio map using this manual survey dataset. Figure 3 illustrates
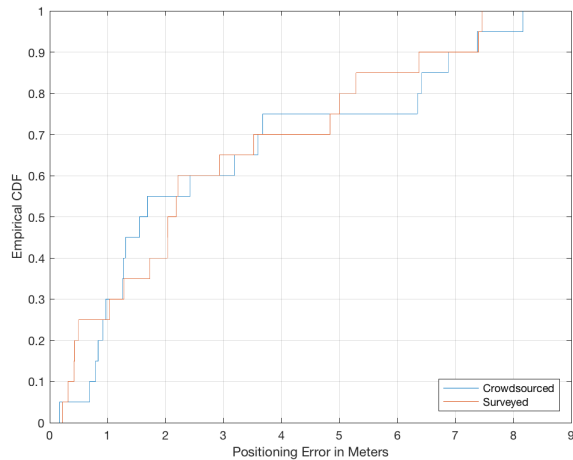
Fig. 4: Empirical Cumulative Distribution of Positioning Error

maps for the same beacons from the crowdsourced and the manual survey.

The second dataset was collected using the same surveying tool, and consisted of twenty fingerprints (ten sampled from the corridors and ten from those ten office rooms) with Tango-derived location tags. We used this dataset to quantitatively evaluate the positioning performance of the maps.

B. Positioning Evaluation

We evaluate the proposed system by comparing the positioning accuracy of the final maps. We adopt the approach in [11] to position the test set based on the crowdsourced and surveyed radio maps respectively. Specifically, a position is determined by maximising the posterior:

$$p(Location|Observation, Radio\,Map)$$

$$\propto p(Location)p(Observation|Location, Radio\,Map)$$

In tracking mode, the positioning can be boosted due to the prior, however for the evaluation purpose, we use a uniform prior to perform one-shot positioning. As a baseline, the surveyed approach reached a mean positioning error across the 20 test sites of 2.86 m. The crowdsourced approach reached a mean positioning error of 2.97 m—see Figure 4 for the empirical cumulative distribution of the accuracy of both approaches and Figure 5 for a visualisation of the errors from both sets of maps.

If we constrained the position to lie only in the corridor, the crowdsourced approach gave a mean error of 2.47 m while the surveyed approach had a mean error of 2.65 m. This is perhaps expected since, with correctly aligned sequences, the crowdsourced data density along the corridor is much higher than that of the manual survey.
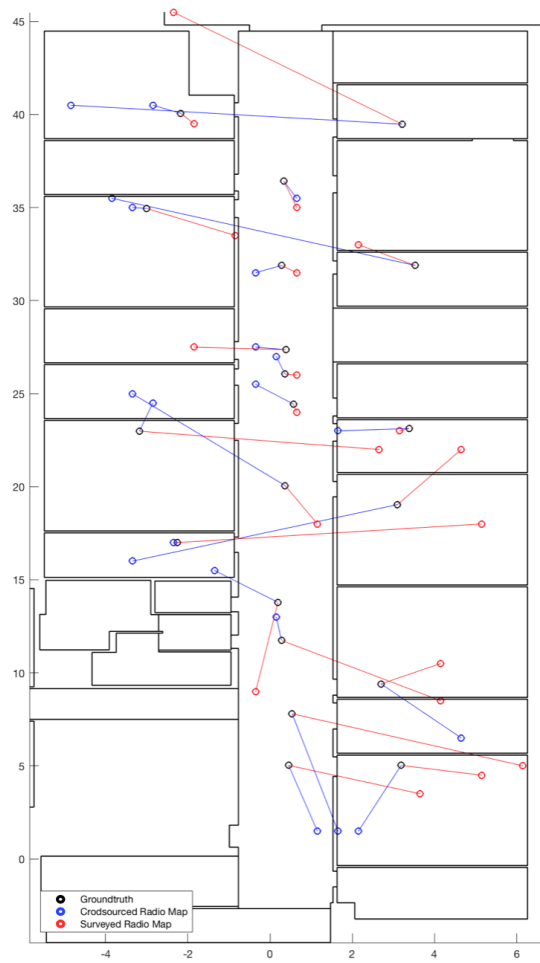

Fig. 5: Positioning Performance of all Fingerprints

V. CONCLUSIONS

We have introduced a promising approach for spatially aligning crowdsourced signal data in order to generate fingerprint maps for indoor positioning. We have demonstrated how to align separate sequences from using regressed signal spaces and then optimising the alignment globally. Our technique produced a radio map for a corridor that was very similar to that from a conventional manual survey, and gave comparable positioning results.

At present our system is based on a number of assumptions about the building. In further work we intend to consider different types of building; the effect of different handsets measuring the signals; the effect of PDR errors; and to test using truly crowdsourced data.

REFERENCES

[1] D. Lymberopoulos, J. Liu, X. Yang, R. R. Choudhury, V. Handziski, and S. Sen, "A realistic evaluation and comparison of indoor location technologies: Experiences and lessons learned," in *Proceedings of the 14th international conference on information processing in sensor networks*. ACM, 2015, pp. 178–189.
[2] D. Han, S. Jung, M. Lee, and G. Yoon, "Building a practical wi-fi-based indoor navigation system," *IEEE Pervasive Computing*, vol. 13, no. 2, pp. 72–79, 2014.

[3] C. Gao and R. Harle, "Sequence-based magnetic loop closures for automated signal surveying," in *Indoor Positioning and Indoor Navigation (IPIN), 2015 International Conference on*. IEEE, 2015, pp. 1–12.

[4] P. Bolliger, "Redpin-adaptive, zero-configuration indoor localization through user collaboration," in *Proceedings of the first ACM international workshop on Mobile entity localization and tracking in GPS-less environments*. ACM, 2008, pp. 55–60.

[5] M. Lee, H. Yang, D. Han, and C. Yu, "Crowdsourced radiomap for room-level place recognition in urban environment," in *Pervasive Computing and Communications Workshops (PERCOM Workshops), 2010 8th IEEE International Conference on*. IEEE, 2010, pp. 648–653.

[6] J.-g. Park, B. Charrow, D. Curtis, J. Battat, E. Minkov, J. Hicks, S. Teller, and J. Ledlie, "Growing an organic indoor location system," in *Proceedings of the 8th international conference on Mobile systems, applications, and services*. ACM, 2010, pp. 271–284.

[7] A. Rai, K. K. Chintalapudi, V. N. Padmanabhan, and R. Sen, "Zee: zero-effort crowdsourcing for indoor localization," in *Proceedings of the 18th annual international conference on Mobile computing and networking*. ACM, 2012, pp. 293–304.

[8] Z. Yang, C. Wu, and Y. Liu, "Locating in fingerprint space: wireless indoor localization with little human intervention," in *Proceedings of the 18th annual international conference on Mobile computing and networking*. ACM, 2012, pp. 269–280.

[9] A. Brajdic and R. Harle, "Walk detection and step counting on unconstrained smartphones," in *Proceedings of the 2013 ACM international joint conference on Pervasive and ubiquitous computing*. ACM, 2013, pp. 225–234.

[10] R. Harle, "A survey of indoor inertial positioning systems for pedestrians," *IEEE Communications Surveys & Tutorials*, vol. 15, no. 3, pp. 1281–1293, 2013.

[11] B. Ferris, D. Haehnel, and D. Fox, "Gaussian processes for signal strength-based location estimation," in *In proc. of robotics science and systems*. Citeseer, 2006.

[12] C. E. Rasmussen, "Gaussian processes for machine learning," 2006.

[13] J. C. Lagarias, J. A. Reeds, M. H. Wright, and P. E. Wright, "Convergence properties of the nelder–mead simplex method in low dimensions," *SIAM Journal on optimization*, vol. 9, no. 1, pp. 112–147, 1998.

[14] J. Biswas, B. Coltin, and M. Veloso, "Corrective gradient refinement for mobile robot localization," in *2011 IEEE/RSJ International Conference on Intelligent Robots and Systems*. IEEE, 2011, pp. 73–78.